

INFORMATION ENTROPY-BASED SCHEDULING FOR COMMUNICATION-EFFICIENT DECENTRALIZED LEARNING

Jaiprakash Nagar^{*}, Zheng Chen[†], Marios Kountouris^{*,‡}, Photios A. Stavrou^{*}

^{*} Communication Systems Department, EURECOM, Sophia-Antipolis, France

[†] Department of Electrical Engineering, Linköping University, Sweden

[‡] DaSCI, Department of Computer Science and Artificial Intelligence, University of Granada, Spain

ABSTRACT

This paper addresses decentralized stochastic gradient descent (D-SGD) over resource-constrained networks by introducing node-based and link-based scheduling strategies to enhance communication efficiency. In each iteration of the D-SGD algorithm, only a few disjoint subsets of nodes or links are probabilistically activated, subject to a given communication cost constraint. We propose a novel importance metric based on information entropy to determine node and link scheduling probabilities. We validate the effectiveness of our approach through extensive simulations, comparing it against state-of-the-art methods, including betweenness centrality (BC) for node scheduling and *MATCHA* for link scheduling. The results show that our method consistently outperforms the BC-based method in node scheduling, achieving faster convergence with up to 60% lower communication budgets. At higher communication budgets (above 60%), our method maintains comparable or superior performance. In link scheduling, our method delivers results that are either superior or on par with those of *MATCHA*.

Index Terms— Decentralized machine learning, information entropy, importance metric, communication network, scheduling.

1. INTRODUCTION

Collaborative training of a machine learning (ML) model with decentralized data and model aggregation offers benefits over centralized learning such as lower communication costs, enhanced privacy, increased scalability, and greater robustness against single point of failure [1, 2]. In such setups, nodes train a shared ML model on their local datasets beyond sharing it with other network nodes (see Fig. 1), thus preserving user privacy [3]. Then, each node shares its locally trained model with the rest of the network nodes by peer-to-peer communication among neighboring nodes, known as decentralized learning.

Decentralized stochastic gradient descent (D-SGD) is a highly used optimization approach for decentralized learning [4]. In every iteration of D-SGD, the nodes compute their stochastic gradient, fuse it with the updates received from the neighbors, and then share their updated model with their neighbors. However, this frequent exchange of model updates results in huge communication overheads, thus making the communication cost a bottleneck. Previously, the convergence speed of D-SGD was analyzed in terms of the decrease in error per iteration [5, 6], neglecting the communication cost of each iteration. However, communication topology significantly affects the convergence speed of D-SGD [7], since the choice of channel access scheme governs the communication cost per iteration. Based on this observation, some works, e.g., [8, 9], used graph

sparsification to enhance convergence speed by reducing communication frequency. Recent works, such as [10], [11], also advocated the use of graph sparsification for communication-efficient decentralized learning as “*not all links are equally important in a graph*”. This means that the most important links activate more often than the least important links to achieve faster convergence. In addition, a node can transmit model updates to all its neighbors in a single transmission slot by exploiting the broadcast nature of communication networks. A more recent work in [12] considers the fact that “*not all nodes are equally important in a graph*”; thus, more critical nodes should communicate more frequently than trivial nodes.

In light of the existing literature, a challenging question is how to measure the importance of links and nodes in the network. In communication networks, each link impacts connectivity differently, having little impact or being crucial. The crucial links for connectivity help to fast error convergence, whereas trivial links can help to reduce communication delays, improving the efficiency of D-SGD [11]. However, the importance of a node in a network reflects its role and impact on the structure and dynamics of the network. Furthermore, computing the importance of the node helps in assessing the vulnerability of the network to node failure and the influence of a node on the flow of information in the network [13]. Additionally, key nodes can affect the convergence speed of decentralized learning algorithms, making it essential to compute link and node importance to estimate their scheduling probabilities for transmitting model updates, thus facilitating fast convergence of the D-SGD algorithm. Although several methods exist for computing link/node importance in a graph, e.g., [14], [15], the work presented in [16] considers the betweenness centrality (BC) method as it renders more accurate node’s ranking in terms of importance than other centrality measures for estimating node activation probabilities. However, the BC method is less effective in identifying critical nodes in large and densely connected topologies [12]. While it offers a global, strategic perspective on key control points within the network, this comes at the expense of high computational complexity and sensitivity to structural changes. Moreover, BC is unsuitable for irregular topologies, as nodes with a degree of one always receive a BC value of zero, regardless of their potential significance.

This paper builds upon the work presented in [16], which employs a BC-based scheduling approach for decentralized learning under partial communication. In contrast, we introduce a novel information entropy (IE)-based metric for assessing the importance of both links and nodes. This metric more accurately identifies critical components of the network than existing centrality measures. By quantifying the information richness and diversity of a node’s local connections, our approach directly captures the structural complex-

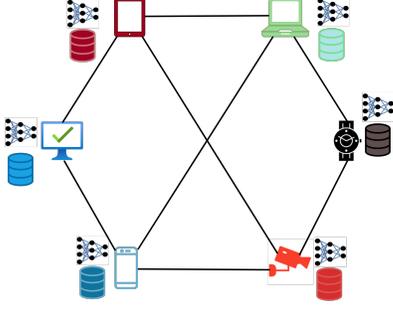


Fig. 1: Distributed machine learning scenario.

ity and the influence of a node's immediate environment. As a result, node importance is evaluated in a manner that is faster, more robust, and more representative of real-world influence, particularly in large, complex, and densely connected networks where local structure is more relevant than global shortest paths. Consequently, the proposed method overcomes the key limitations of traditional centrality-based approaches, enabling efficient and scalable computation of node and link importance in both sparse and dense, irregular topologies. In addition, the proposed approach avoids monotonic importance ranking, enabling more accurate learning with fewer transmission slots. Experimental results show that the proposed IE-based link importance metric consistently performs as well as or better than the optimized MATCHA framework. Furthermore, the IE-based node importance metric significantly outperforms the BC-based approach, achieving higher test accuracy and lower training loss for the D-SGD algorithm while using fewer communication rounds. Simulation studies demonstrate performance improvements of up to 35%, particularly when up to 60% of the node subsets are activated, across the network topologies evaluated in this work.

Notation: We consider a connected network consisting of N nodes, represented by an undirected graph $G = (V, E)$, referred to as the base topology. The node set is denoted by $V = \{1, \dots, N\}$, and the edge set $E \subseteq \{(i, j) \in V \times V | i \neq j\}$ specifies which pair of nodes can directly communicate. For each node $i \in V$, its set of neighbors is defined as $\mathcal{N}_i = \{j | (i, j) \in E\}$, and its degree is given by $d_i = |\mathcal{N}_i|$. The topology of the graph G can be described by its adjacency matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$, where $A_{ij} = 1$ if $(i, j) \in E$, and $A_{ij} = 0$ otherwise. The Laplacian matrix \mathbf{L} associated with the graph is defined as $\mathbf{L} := \mathbf{D} - \mathbf{A}$, where \mathbf{D} is a diagonal matrix with entries $D_{ii} = d_i$, representing the degree of node i .

2. DECENTRALIZED MACHINE LEARNING

Consider a network consisting of N nodes, where each node $i \in V$ has access to a local training dataset \mathcal{D}_i . The primary objective is for all nodes to collaboratively train an ML model by sharing model parameters with their immediate neighbors. Specifically, each node aims to estimate a common parameter vector $\mathbf{x} \in \mathbb{R}^d$ that minimizes the following global objective function: $F(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N F_i(\mathbf{x})$, where $F_i(\mathbf{x})$ denotes the local objective function at node i , defined as $F_i(\mathbf{x}) := \frac{1}{|\mathcal{D}_i|} \sum_{s \in \mathcal{D}_i} \ell(\mathbf{x}; s)$, and $\ell(\mathbf{x}; s)$ is the loss function evaluated at sample s from the local dataset \mathcal{D}_i .

We adopt a consensus-based D-SGD algorithm for collaborative model training, where each node i computes local gradients using its own dataset \mathcal{D}_i , and performs consensus-based model averaging

with its neighbors [15]. Specifically, each iteration of the algorithm consists of the following steps:

- **Stochastic gradient computation:** At iteration k , node i computes a stochastic gradient vector $\mathbf{g}_i^{(k)} = \frac{1}{|\zeta|} \sum_{s \in \zeta} \nabla \ell(\mathbf{x}_i^{(k)}; s)$, where $\zeta \subseteq \mathcal{D}_i$ is a randomly sampled mini-batch from the local dataset \mathcal{D}_i . The node then performs a local update of its model parameters using the stochastic gradient: $x_i^{(k+\frac{1}{2})} = x_i^{(k)} - \gamma g_i^{(k)}$, where γ is the learning rate.
- **Communication with neighbors:** Each node exchanges model parameters with its immediate neighbors to facilitate consensus. Specifically, node i first transmits its updated model parameters $\mathbf{x}_i^{(k+\frac{1}{2})}$ to its neighbors, and then receives the corresponding model parameters from them, denoted as \mathcal{N}_i . As a result, node i obtains both its own and its neighbors' intermediate model parameters for the current iteration.
- **Consensus-based model averaging:** Each node i updates its model parameters by averaging its own and its neighbors' intermediate updates using $x_i^{(k+1)} = \sum_{j=1}^N W_{ij}^{(k)} x_j^{(k+\frac{1}{2})}$, where $W_{ij}^{(k)}$ denotes the weight assigned by node i to the model parameters received from node j during iteration k . These weights can be collectively written in the form of a **weight** or a **mixing matrix** $\mathbf{W}^{(k)} \in \mathbb{R}^{N \times N}$. Note that $W_{ij} \neq 0$ only if $(i, j) \in E$, i.e., if node j is a neighbor of node i .

2.1. Mixing Matrix Properties and Design

According to [17], the convergence of the D-SGD algorithm is guaranteed under the following sufficient conditions:

- the mixing matrix \mathbf{W} is symmetric and doubly stochastic;
- the spectral gap of \mathbf{W} is strictly positive.

In addition to the aforementioned conditions on \mathbf{W} , the convergence of the D-SGD algorithm also requires that the variances of the stochastic gradients \mathbf{g}_i are bounded and that the local objective functions $F_i(\mathbf{x})$ are differentiable with Lipschitz gradients [12]. A common and effective design for the mixing matrix is based on the graph Laplacian matrix [11], and is given by $\mathbf{W} := \mathbf{I} - \alpha \mathbf{L}$, where \mathbf{I} is the identity matrix, $\alpha > 0$ is a tunable parameter, and \mathbf{L} is the Laplacian matrix of the communication graph. The parameter $\alpha > 0$ should be chosen to ensure that the spectral radius satisfies: $\rho(\mathbf{W} - \mathbf{J}) < 1$, where $\rho(\cdot)$ denotes the spectral radius of a matrix, $\mathbf{J} = \frac{1}{N} \mathbf{1}\mathbf{1}^T$, and $\mathbf{1}$ is the all-ones column vector. Furthermore, for faster convergence, α should be selected to minimize $\rho(\mathbf{W} - \mathbf{J})$.

3. COMMUNICATION-EFFICIENT DECENTRALIZED LEARNING UNDER PARTIAL COMMUNICATION

In this section, we describe the communication model, introduce the proposed scheduling schemes based on node importance metrics, and explore strategies for optimizing the mixing matrices. To ensure convergence in distributed learning, nodes must exchange model parameters with their neighbors based on the communication model. Due to the broadcast nature of wireless channels, a node can share its updated model with all neighbors in a single time slot. However, to prevent interference and packet collisions from simultaneous transmissions, nodes must coordinate their communication schedules with neighboring nodes, ensuring reliable communication and algorithmic convergence. The most effective way

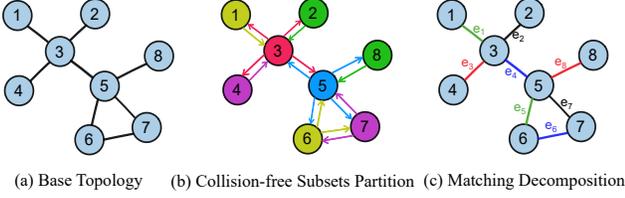


Fig. 2: Partitioning (a) Undirected base topology into (b) Collision-free Subsets and (c) Matchings.

to ensure collision-free communication is by allocating orthogonal communication resources, such as distinct time slots, to each broadcasting node. In this scheme, every node is assigned a dedicated time slot during which it can broadcast its model parameters to its neighbors without interference. Consequently, a complete communication round, in which all nodes exchange updates, requires a total of N time slots to ensure full information fusion across the network. Therefore, under a full communication scenario, executing \mathcal{K} iterations of the consensus algorithm necessitates a total of $N\mathcal{K}$ time slots.

3.1. Communication Model

We consider two types of communication scenarios: (1) peer-to-peer communication with link-based scheduling, and (2) broadcast transmission with node-based scheduling. In the first scenario, links are activated randomly, and each information exchange over a link requires two transmission slots. In the second, nodes are activated randomly, and each broadcast allows a node's model update to reach multiple neighbors simultaneously. In both cases, we assume that simultaneous transmissions adhere to collision-free conditions.

3.1.1. Link Scheduling with Peer-to-Peer Transmission

A matching is a subgraph of G where each node is connected to at most one edge. The base topology is decomposed into M disjoint matching subgraphs $\{G_j(V, E_j)\}_{j=1}^M$, where $E = \bigcup_{j=1}^M E_j$ and $E_i \cap E_j = \emptyset, \forall i \neq j$. Fig. 2(c) shows the matching decomposition of a base topology using the edge-coloring algorithm [18], where edges with the same color form a matching. Note that in each communication round, only a subset of these matchings is activated according to the scheduling policy. This scheduling is guided by the principle that links critical for maintaining network connectivity should be activated more frequently than less essential ones.

3.1.2. Node Scheduling with Broadcast

A collision-free subset is a group of nodes where no two nodes share a common neighbor, enabling them to transmit model updates simultaneously without the risk of collisions. The base topology is partitioned into a finite number of such collision-free subsets, denoted by q . Fig. 2(b) illustrates the partitioning of the base topology, where nodes with the same color form a collision-free subset that requires only a single transmission slot. Each communication round of the D-SGD algorithm is further divided into multiple transmission time slots. For a given node pair (i, j) , node j can successfully receive information from node i only if neither node j nor any of its neighbors are transmitting during that slot.

An undirected graph $G = (V, E)$ can be represented by a directed graph $G^d = (V, E^d)$, where each undirected edge is replaced by a pair of bidirectional directed links connecting each pair

of nodes. Due to the broadcast nature of communication channels, each node transmits its information to all of its neighbors via directed links, resulting in a local star topology. To enable collision-free concurrent transmissions, the base topology is partitioned into disjoint subsets of nodes $\{\mathcal{V}_r\}_{r \in [q]}$, where $q \leq N$. Each subset \mathcal{V}_r contains nodes that do not share any common neighbors, allowing them to broadcast simultaneously without interference, $i, j \in \mathcal{V}_r$ if $\mathcal{N}_i \cap \mathcal{N}_j = \emptyset$. These subsets form a complete partition of the node set, satisfying $\bigcup_{r=1}^q \mathcal{V}_r = \mathcal{V}$, $\mathcal{V}_r \cap \mathcal{V}_l = \emptyset, \forall r \neq l$.

We employed a greedy vertex-coloring algorithm [19] to partition the base topology by assigning different colors to connected nodes. To enforce the constraint that no two nodes sharing a common neighbor can broadcast concurrently, even if they belong to the same subset, we construct an auxiliary graph $G^a = (V, E^a)$, where E^a includes all original edges from E along with additional edges connecting every pair of nodes that share at least one common neighbor. This ensures that nodes with the same color, and thus in the same subset, form a collision-free set. The total number of such subsets, denoted by q , corresponds to the number of colors assigned by the coloring algorithm.

Furthermore, this work considers a partial communication policy aimed at reducing the number of transmission slots per iteration and minimizing the overall communication cost required to achieve consensus in a distributed learning algorithm. Under partial communication, only a subset of network links and nodes are activated for information exchange in each round (see, e.g., [16]), implying that communication should not be uniformly distributed across all links and nodes. Instead, the frequency of communication should be guided by the relative importance of individual links and nodes within the graph. To this end, this paper employs an information entropy-based approach to quantify the importance of links and nodes, as detailed in Section 3.2.

3.2. Node Importance Vector

In what follows, we describe the IE-based node importance approach, which leverages node degree to compute the self-information (SI) [20] of an edge connecting two adjacent nodes [21]. Specifically, the method quantifies the informational contribution of a link $(i, j) \in E$, using the degrees of the two nodes it connects, d_i and d_j . The probability associated with the existence of link (i, j) is defined as $P(i, j) = \frac{1}{d_i d_j}$. Consequently, the self-information of the link, interpreted as its weight, is computed as:

$$SI(i, j) = -\log_2 P(i, j) = \log_2(d_i d_j).$$

For undirected links, the value of $SI(i, j) = SI(j, i)$. We now define two notations $S(i)$ and $S^+(i)$, which are important to calculate the node importance score. $S(i)$ denotes the sum of SI values of all links connected to node i and is calculated by $S(i) = \sum_{j \in \mathcal{N}_i} SI(i, j)$. $S^+(i)$ captures the combined structural importance of links incident to node i and its neighbors, defined as the sum of SI of links that node i and its neighbors are one endpoint of these links, i.e., $S^+(i) = \sum_{j \in \mathcal{N}_i \cup \{i\}} S(j)$. Note that $S(i)$ accounts for the influence of links directly connected to node i , while $S^+(i)$ reflects a broader view by incorporating the link contributions of node i and its neighbors. The probability assigned to any node $j \in \mathcal{N}_i \cup \{i\}$, denoted by $P(j)$, can be defined as

$$P(j) := \frac{S(j)}{S^+(i)}, \quad (1)$$

such that the sum of all the probabilities is equal to 1, i.e.,

$\sum_{j \in \mathcal{N}_i \cup \{i\}} \frac{S(j)}{S^+(i)} = 1$, and the IE of node i is calculated as $\mathcal{E}(i) = -\sum_{j \in \mathcal{N}_i \cup \{i\}} P(j) \log_2 P(j)$.

The amount of information a node possesses is influenced by the number and strength of its connections within the graph. To quantify this, we compute the expected entropy of node i , denoted by $\mathcal{E}(i)$, which incorporates the IE of neighboring nodes, weighted by link importance as a measure of relevance. Nodes are then ranked based on their expected entropy values: a higher $\mathcal{E}(i)$ indicates a more important node. In cases where multiple nodes share the same expected entropy, they are assigned the same rank. For example, in a network of six nodes, if nodes 1 and 2 have equal expected entropy, and nodes 4 and 5 share another common value, the ranking would be: $\mathcal{E}(1) = \mathcal{E}(2) > \mathcal{E}(3) > \mathcal{E}(4) = \mathcal{E}(5) > \mathcal{E}(6)$. To assess the importance of links within the graph, we transform the original graph into its corresponding line graph, where each node in the line graph represents an edge from the original graph [22]. We then apply the IE-based method to this line graph. The resulting node importance vector in the line graph directly corresponds to the link importance in the original graph.

3.3. Probabilistic Scheduling Design

Similar to [11, 12], we consider a probabilistic scheduling method, where an arbitrary number of collision-free subsets are activated in each communication round. The activation decision of a matching in each communication round of D-SGD algorithm is an independent Bernoulli random variable Z_j , which is equal to 1 with probability p_j and 0 otherwise, for each matching $G_j, \forall j \in \{1, \dots, M\}$. Here, p_j represents the expected activation probability of the j^{th} matching.

Further, for node-based scheduling, each collision-free subset has a certain scheduling probability. Let the scheduling decision of collision-free subset S_r in k -th communication round be $s_i(k)$, where $s_i = 1$ if S_r is activated, and 0 otherwise. The graph of the communication topology activated in communication round k is represented as $G(k) = \{V(k), E(k)\}$, where $G(k) = \bigcup_{i=1}^q s_i(k) S_i$. Note that $G(k) = G^d$ when all the collision-free subsets are activated in communication round k . Note that the sum of the activation probabilities of all matchings (for link-based scheduling) and collision-free subsets (for node-based scheduling) should be less than or equal to the communication budget, denoted by B .

3.4. Optimizing Scheduling Probabilities and Mixing Matrix

We can get the topology of the scheduled broadcasting nodes by sampling the corresponding columns of the adjacency matrix of the base topology, leaving the columns unchanged, and masking the remaining columns as zero. The obtained matrix will represent a directed graph with edges connecting activated nodes (columns) to their neighbors. A network with bidirectional links can be obtained by sampling the columns and rows with the same indices (i.e., i -th column and row) from the adjacency matrix. The resultant adjacency and Laplacian matrices will be symmetric, and as a result, the mixing matrix of the scheduled broadcasting nodes will also be symmetric.

Mathematically, a symmetric adjacency matrix $\hat{\mathbf{A}}$ can be obtained by sampling the columns and rows of \mathbf{A} , as $\hat{\mathbf{A}}(k) = \mathbf{Q}(k) \mathbf{A} \mathbf{Q}(k)$, where $\mathbf{Q}(k) = \text{diag}(n_i(k), \dots, n_N(k))$, with $n_i(k) = 1$ if node i is activated in k -th iteration. It is cardinal to note that scheduling decisions are made only on collision-free subsets, not on individual nodes, and multiple collision-free subsets can be scheduled in each communication round. The Laplacian matrix used for the scheduled communication topology in k -th communication

round is $\hat{\mathbf{L}}(k) = \text{diag}(\hat{\mathbf{A}}(k) \mathbf{1}) - \hat{\mathbf{A}}(k)$, where the resulting Laplacian matrix $\hat{\mathbf{L}}(k)$ is symmetric, and the sum of each row/column is zero. Thus, the resultant mixing matrix for k -th communication round is also symmetric, defined as $\mathbf{W}(k) = \mathbf{I} - \alpha \hat{\mathbf{L}}(k)$, where α is a constant independent of k .

The scheduling probability of a networked node depends on its importance in the network topology, where an important node would be scheduled more often as compared to a trivial networked node. In every communication round, a subset S_r is scheduled with probability p_{S_r} , for $r \in \{1, \dots, q\}$, thus $\mathbb{E}[\mathbf{Q}(k)] = \text{diag}(p_1, \dots, p_N)$. Note that a node's activation probability is equal to the scheduling probability of the subset to which it belongs, i.e., $p_i = p_{S_r}$ if $i \in \mathcal{V}_r$. As a result, the number of subsets scheduled in a communication round is an arbitrary number such that the average required count of transmission slots per communication round is $\sum_{l=1}^q p_{S_l}$.

The subset scheduling probabilities had been optimized in [16] to accelerate per-round convergence under constrained communication as follows:

$$\begin{aligned} & \min_{p_{S_1}, \dots, p_{S_q}} \|\mathbb{E}[\mathbf{W}^2(t)] - \mathbf{J}\|_2, \\ \text{s.t. } & \sum_{l=1}^q p_{S_l} = B \text{ and } 0 \leq p_{S_l} \leq 1, \forall l. \end{aligned} \quad (2)$$

In (2), B denotes the average count of transmission slots per communication round. For a detailed explanation of the optimization problem (2), the interested reader is referred to [16]. In this work, we use the IE-based approach to rank networked nodes and compute their activation probabilities. Here, the activation probability of node i is represented by b_i such that $\sum_{i=1}^N b_i = 1$, and the scheduling probability of a subset j is given by $b_{S_j} = \sum_{i=1}^N b_i \mathbf{1}_{\{i \in \mathcal{V}_j\}}$, where the value of $\mathbf{1}_{\{i \in \mathcal{V}_j\}}$ is equal to one if $i \in \mathcal{V}_j$, and zero otherwise. We also choose $p_{S_j} = \min\{1, \gamma b_{S_j}\}$, where the value of γ is computed such that $\sum_{i=1}^q p_{S_i} = B$.

For every communication round, it is pivotal to design the mixing matrix as $\mathbf{W}(k) = \mathbf{I} - \alpha \hat{\mathbf{L}}(k)$ to enhance the convergence speed of the D-SGD. Moreover, the mixing matrix can be optimized by solving (2) considering α as an optimization parameter. The solution to this convex problem is provided in [16] and is reproduced here for the sake of completeness as follows:

$$\begin{aligned} & \min_{s, \alpha, \beta} s, \\ \text{s.t. } & \alpha^2 - \beta \leq 0, \\ & \mathbf{I} - 2\alpha \mathbb{E}[\hat{\mathbf{L}}(k)] + \beta \left(\mathbb{E}[\hat{\mathbf{L}}^T(k) \hat{\mathbf{L}}(k)] \right) - \mathbf{J} \leq s \mathbf{I}. \end{aligned} \quad (3)$$

where s and β are auxiliary variables.

4. EXPERIMENTAL RESULTS

In this section, we evaluate the performance of the IE-based link/node importance metric in terms of test accuracy and training loss within the D-SGD framework, following the experimental setup described in [16]. We report results for three different network topologies comprising $N = \{15, 20, 30\}$ nodes, as shown in Fig. 3. As illustrated in Fig. 3, we compare the performance of IE-based and BC-based node importance metrics in three different network topologies. The percentage shown in each figure indicates the average fraction of activated subsets per iteration. While we experimented with various activation percentages across all topologies,

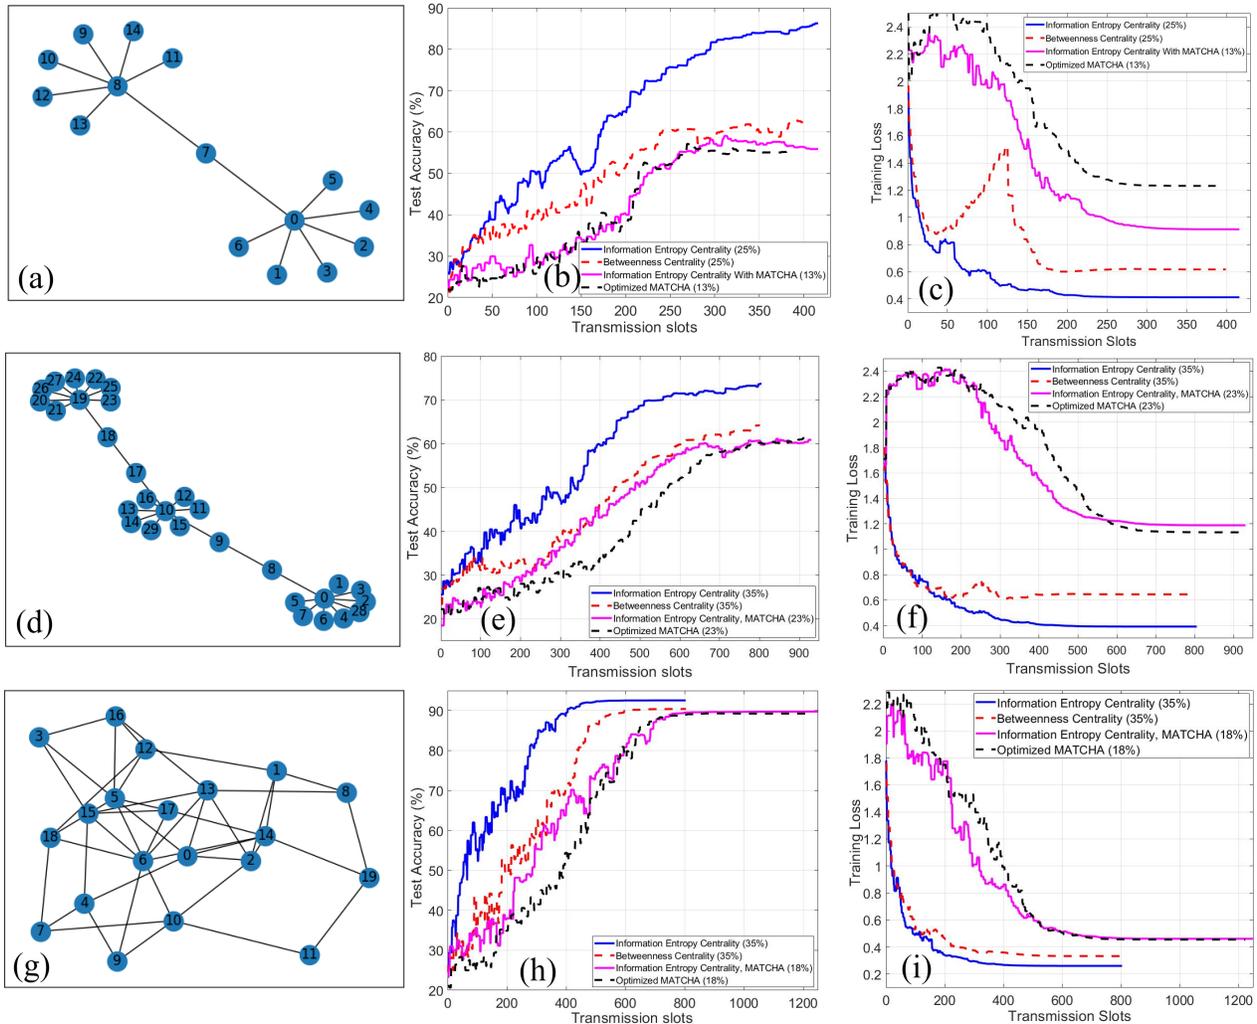


Fig. 3: Performance comparison between IE and BC-based link/node importance methods.

we report results only for lower activation levels due to space limitations. The IE-based node importance metric consistently achieves higher test accuracy in fewer transmission slots than the BC-based method when up to 60% of subsets are activated in all three network topologies. For instance, in Fig. 3(b), with 25% subset activation in the 2-star topology, the IE-based approach reaches 70% test accuracy in approximately 200 transmission slots, whereas the BC-based method fails to achieve 70% accuracy even after 400 slots. A similar pattern is observed in the 3-star topology at 35% subset activation: the IE-based approach achieves 70% accuracy in roughly 500 slots, while the BC method remains below this threshold even after 800 slots, as shown in Fig. 3(e). We observe similar trends for random topologies. At 35% activation, the IE-based method reaches 90% test accuracy in around 400 transmission slots, whereas the BC-based approach requires nearly 600 slots to attain the same accuracy, as illustrated in Fig. 3(h). Furthermore, the IE-based link importance metric achieves target accuracy using fewer or equal transmission slots compared to the optimized MATCHA approach, as demonstrated in Figs. 3(b), (e), and (h). For example, in the 3-star topology with 23% activated matchings, the IE-based method

attains 60% test accuracy in roughly 650 transmission slots, while optimized MATCHA requires approximately 770 slots. Conversely, both methods reach about 90% test accuracy at nearly 800 slots (Fig. 3(e)). Importantly, the training loss associated with the IE-based method remains consistently lower than that of both the BC approach and optimized MATCHA, apart from some initial fluctuations, as shown in Figs. 3(c), (f), and (i). Overall, the proposed IE-based method outperforms the BC approach in large, dense, and sparse networks. Its advantage lies in capturing the richness of local information and the structural complexity of the network, enabling the identification of influential nodes based on their diverse neighborhood structures. This leads to significantly higher test accuracy with fewer transmission slots compared to the BC-based method.

5. CONCLUSION

In this work, we proposed an IE-based link and node importance metric to compute the scheduling probabilities of disjoint subsets of links or nodes, aiming to accelerate decentralized learning under partial communication constraints. The key novelty of the IE-based approach lies in its dual use of self-information-weighted links

and neighborhood entropy to quantify node importance, embedding structural complexity into the ranking process. This makes the approach scalable and robust, thus outperforming traditional path- and degree-based metrics. Simulation results demonstrate that the IE-based approach outperforms the BC method, achieving faster convergence with fewer transmission slots under lower communication budgets (up to 60%) across various network topologies, including dense random graphs and sparse, irregular structures. At higher communication budgets (above 60%), the IE-based method delivers performance that is comparable to or better than the BC method across all tested topologies. Additionally, for link-based scheduling scenarios, the proposed method achieves performance that is either superior to or on par with *MATCHA*. A promising direction for future research involves developing node scheduling strategies that jointly consider both the volume of data at each node and the distributions of local datasets.

6. REFERENCES

- [1] Shuyan Hu, Xiaojing Chen, Wei Ni, Ekram Hossain, and Xin Wang, "Distributed machine learning for wireless communication networks: Techniques, architectures, and applications," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 3, pp. 1458–1493, 2021.
- [2] Enrique Tomás Martínez Beltrán, Mario Quiles Pérez, Pedro Miguel Sánchez Sánchez, Sergio López Bernal, G r me Bovet, Manuel Gil P rez, Gregorio Mart nez P rez, and Alberto Huertas Celdr n, "Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges," *IEEE Communications Surveys & Tutorials*, 2023.
- [3] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [4] Brian Swenson, Ryan Murray, H Vincent Poor, and Soumya Kar, "Distributed stochastic gradient descent: Nonconvexity, nonsmoothness, and convergence to local minima," *Journal of Machine Learning Research*, vol. 23, no. 328, pp. 1–62, 2022.
- [5] Jianyu Wang and Gauri Joshi, "Cooperative SGD: A unified framework for the design and analysis of local-update sgd algorithms," *Journal of Machine Learning Research*, vol. 22, no. 213, pp. 1–50, 2021.
- [6] D. Jakovetic, D. Bajovic, A. K. Sahu, and S. Kar, "Convergence rates for distributed stochastic optimization over random networks," in *CDC 2018-IEEE Conference on Decision and Control*. 2018, pp. 4238–4245, IEEE.
- [7] Angelia Nedi c, Alex Olshevsky, and Michael G Rabbat, "Network topology and communication-computation tradeoffs in decentralized optimization," *Proceedings of the IEEE*, vol. 106, no. 5, pp. 953–976, 2018.
- [8] Jianyu Wang and Gauri Joshi, "Adaptive communication strategies to achieve the best error-runtime trade-off in local-update sgd," *Proceedings of Machine Learning and Systems*, vol. 1, pp. 212–229, 2019.
- [9] Nguyen H Tran, Wei Bao, Albert Zomaya, Minh NH Nguyen, and Choong Seon Hong, "Federated learning over wireless networks: Optimization model design and analysis," in *INFOCOM 2019-IEEE conference on computer communications*. IEEE, 2019, pp. 1387–1395.
- [10] Cho-Chun Chiu, Xusheng Zhang, Ting He, Shiqiang Wang, and Ananthram Swami, "Laplacian matrix sampling for communication-efficient decentralized learning," *IEEE Journal on Selected Areas in Communications*, vol. 41, no. 4, pp. 887–901, 2023.
- [11] Jianyu Wang, Anit Kumar Sahu, Gauri Joshi, and Soumya Kar, "Matcha: A matching-based link scheduling strategy to speed up distributed optimization," *IEEE Transactions on Signal Processing*, vol. 70, pp. 5208–5221, 2022.
- [12] Daniel P rez Herrera, Zheng Chen, and Erik G Larsson, "Faster convergence with less communication: Broadcast-based subgraph sampling for decentralized learning over wireless networks," *arXiv preprint arXiv:2401.13779*, 2024.
- [13] Rongrong Yin, Xueliang Yin, Mengdi Cui, and Yinghan Xu, "Node importance evaluation method based on multi-attribute decision-making model in wireless sensor networks," *EURASIP Journal on Wireless Communications and Networking*, vol. 2019, pp. 1–14, 2019.
- [14] Yuanzhi Yang, Lei Yu, Zhongliang Zhou, You Chen, and Tian Kou, "Node importance ranking in complex networks based on multicriteria decision making," *Mathematical Problems in Engineering*, vol. 2019, no. 1, pp. 9728742, 2019.
- [15] Daniel P rez Herrera, Zheng Chen, and Erik G Larsson, "Distributed consensus in wireless networks with probabilistic broadcast scheduling," *IEEE Signal Processing Letters*, vol. 30, pp. 41–45, 2023.
- [16] Daniel P rez Herrera, Zheng Chen, and Erik G Larsson, "Decentralized learning over wireless networks with broadcast-based subgraph sampling," in *ICC 2024-IEEE International Conference on Communications*. IEEE, 2024, pp. 932–937.
- [17] Hao Ye, Le Liang, and Geoffrey Ye Li, "Decentralized federated learning with unreliable communications," *IEEE journal of selected topics in signal processing*, vol. 16, no. 3, pp. 487–500, 2022.
- [18] Jayadev Misra and David Gries, "A constructive proof of vizing's theorem," *Information Processing Letters*, vol. 41, no. 3, pp. 131–133, 1992.
- [19] B la Bollob s, *Modern graph theory*, vol. 184, Springer Science & Business Media, 1998.
- [20] Claude Elwood Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, vol. 27, pp. 379–423, 1948.
- [21] Shihu Liu and Haiyan Gao, "The self-information weighting-based node importance ranking method for graph data," *Entropy*, vol. 24, no. 10, pp. 1471, 2022.
- [22] Omar De la Cruz Cabrera, Mona Matar, and Lothar Reichel, "Edge importance in a network via line graphs and the matrix exponential," *Numerical Algorithms*, vol. 83, pp. 807–832, 2020.